DOCUMENT RESUME

ED 282 521 IR 012 664

AUTHOR Ohlsson, Stellan
TITLE Computer Simulation and Its Impact on Educational Research and Practice.
INSTITUTION Pittsburgh Univ., Pa. Learning Research and Development Center.
SPONS AGENCY National Science Foundation, Washington, D.C.; Pittsburgh Univ., Pa. Learning Research and Development Center.
PUB DATE 86
GRANT MDR-8470339; OERI-G-86-0005
NOTE 65p.
PUB TYPE Information Analyses (070) -- Reports - Research/Technical (143)

EDRS PRICE MF01/PC03 Plus Postage.
DESCRIPTORS Cognitive Measurement; Cognitive Processes; *Cognitive Psychology; *Computer Assisted Instruction; *Computer Simulation; *Educational Research; Futures (of Society); Learning Theories; Literature Reviews; Models; *Psychological Studies; Research Methodology
IDENTIFIERS *Connectionist Approach

ABSTRACT

This essay summarizes the theory and practice of computer simulation, assesses the state of the art of simulation with respect to pedagogically relevant processes like learning, and speculates about the impact of such simulations on pedagogical research and practice. Arguing that the use of computer simulation as a technique for building formal models of mental processes forces the cognitive psychologist to consider the content of strategic or heuristic knowledge, the paper begins by discussing such philosophical concepts as the formalization of theories, the distinction between theories and models, and the notion of a research program. The rationale and workmode of simulation research are then summarized, and a review of the literature illustrates the range of phenomena with educational relevance to which the simulation technique has been applied. The new connectionist approach to simulation is described, and concern is expressed about the way in which knowledge appears in connectionist theories. The most direct interaction of computer simulation with education in the future is predicted to be through such computerized teaching tools as intelligent tutoring systems and systems for automatic cognitive diagnosis. It is concluded that a traditional science-to-practice knowledge transfer will occur to the extent that simulation models contribute to the improvement of psychological theories with pedagogical relevance, and that computerized teaching devices will have a dramatic effect on cognitive research methodology by providing access to information on the behavior of students in real learning situations. A list of 170 references is provided. (MES)

LEARNING RESEARCH AND DEVELOPMENT CENTER

1986/14

COMPUTER SIMULATION AND ITS IMPACT ON EDUCATIONAL RESEARCH AND PRACTICE

STELLAN OHLSSON

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☒ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

J. Aug

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

University of Pittsburgh

BEST COPY AVAILABLE

ED282521

IR012664

COMPUTER SIMULATION AND ITS IMPACT ON EDUCATIONAL RESEARCH AND PRACTICE

Stellan Ohlsson

Learning Research and Development Center
University of Pittsburgh

1986

To appear in *International Journal of Education*, in press.

The research reported herein was conducted at the Learning Research and Development Center, funded in part by the National Institute of Education (NIE), U. S. Department of Education. Preparation of the manuscript was supported in part by an NSF grant and by a grant for the OERI Center for the Study of Learning. The opinions expressed do not necessarily reflect the position or policy of the sponsoring agencies, and no official endorsement should be inferred.

Copyright © 1986 Stellan Ohlsson

Abstract

The use of computer simulation as a technique for building formal models of mental processes forces the cognitive psychologist to consider the content of knowledge, in particular strategic or heuristic knowledge. The rationale and workmode of simulation research are summarized. A short review illustrates the range of phenomena with educational relevance to which the simulation technique has been applied. Computer simulation and education are predicted to interact in several ways in the future, most directly through computerized teaching tools like intelligent tutoring systems and systems for automatic cognitive diagnosis.

> ***In the practical business of giving explanations--after all--scientists often enough rely not on the presentation of deductive arguments, but on such alternative activities as the drawing of graphs or ray-diagrams, the construction of intellectual models, or the programming of computers.... [The] primary thing to be learned, tested, put to work, criticized, and changed [is] the repertory of intellectual techniques, procedures, skills, and methods of representation, which are employed in 'giving explanations' of events and phenomena within the scope of the science concerned.***
>
> Stephen Toulmin, *Human Understanding*, pp. 157-159.

Teachers, I am sure, strive to teach according to how learners learn. Their ambition is hampered, however, by our lack of knowledge about the psychology of learning. In the past, pedagogical applications have been derived from such disparate ideas as the behaviorist concept of reinforcement (Glaser, 1978) and the piagetian notion of stages of intellectual growth (Groen, 1978), but they have been largely unsuccessful[1]. The art of teaching has been unable to turn itself into the engineering of instruction, for lack of a viable science of learning.

In the last three decades, the information processing perspective has become the dominant point of view in the study of thinking, reasoning, problem solving, memory, learning, attention, perception, and other knowledge-related mental processes and functions (see, e.g., Anderson, 1980; Gardner, 1985). The central notion of this new psychology of cognition is that knowledge resides in symbol structures in the head, and that mental processes consist of using, editing, or creating those symbol structures. The main goals of this essay are to summarize the theory and practice of computer simulation, to assess the state of the art of simulation with respect to pedagogically relevant processes like learning, and to speculate about the impact of such simulations on pedagogical research and practice.

Even though the psychology of mental symbol manipulations is a young creature, it now has

---

[1] See Good & Brophy (1986) for an up to date introduction to educational psychology.

offspring. One group of cognitive psychologists is already changing the way we think about computer simulation by basing simulations on what we know about how the brain works, as well as on analyses of knowledge. This research orientation - often labeled "connectionism" - uses the research tool of computer simulation within a new theoretical framework, producing simulations of a novel sort (McClelland, Rumelhart, & the PDP research group, in press a,b). This is an interesting and important development. The focus of this essay is nevertheless on what we now must call traditional computer simulations, while connectionist simulations will be discussed only briefly. The main reason for this bias is that the application of connectionist research to education is, by and large, yet to come, while computer simulationists in general have payed considerable attention to classroom tasks in recent years.

The purpose of the first section below is to remind the reader of some well-known but useful philosophical concepts, such as the idea of formalization of theories, the distinction between theories and models, and the notion of a research programme. The second and main section explains, reviews, and critiques the computer simulation programme, concluding that there are difficulties with it, some - but only some - of which could be overcome if researchers took the programme seriously. In a short section on the new programme of connectionism I describe its methodological advantages and express my concern about the way in which knowledge appears (or, rather, does not appear) within connectionist theories. I then turn to the question of what the pedagogical consequences of the computer simulation technique might be, reaching the conclusion that although simulation lends itself to pedagogical purposes, the impact of simulation research on educational practice will be slight. Instead, computer simulation makes it possible for educational practice to set the agenda for cognitive research.

On Theories and Theorizing

Theorizing is a complex activity, and theories have many aspects; let us illustrate them with a handful of metaphors. A scientific theory represents an effort at understanding. Its purpose is to

enable clear, coherent, and successful thinking about some part or aspect of the world. A theory is thus like a weave of ideas, its strength a function of how tight the strands are woven together. Penetrating behind the observable surface of the world, a theory postulates - creates, invents - objects, processes, interactions, etc., such that if they were real, the surface characteristics of the world would have to be just the way we observe them to be. A theory is thus like a fantasy or a daydream, albeit a very disciplined and purposeful one. If the world was described in a long and detailed text, a theory would be a statement of the main ideas, the themes which underly the text as a whole. A theory is thus like an abstract. Once proposed, a successful theory is typically applied again and again to one phenomenon after another. If it continues to bring clarity and understanding, more and more phenomena will be seen as belonging to its domain. A theory is thus like a politically ambitious emperor, reaching out to conquer a larger and larger territory. The effort to extend a theory leads to the clarification of concepts, the sharpening of distinctions, and the exploration of alternative formulations and formalizations. A theory is thus like a crystal, steadily growing into its perfect form.

Since theories deal with entities which cannot be observed directly, theoreticians have a unique problem in communicating with each other: They must dispense with what philosophers call "ostensive definitions", i. e., they cannot point to instances of what they are talking about; they cannot explain what they mean by showing examples. Since the sensory content of scientific terms is thus minimized, the rules by which those terms are combined becomes so much more important. Scientific discourse takes on the character of a language game, the central terms becoming mere tokens fitted into standardized patterns of expression.

It is not only hard to communicate about unobservables, it is also hard to think about them. Ordinarily, we think within the context of the common sense knowledge about our world that we have collected by being participants in it. But none of us have gamboled with electrons and quarks as children, conversed at length with double helixes, or spent Sunday afternoons

observing patterns of activations spread across semantic networks. Reasoning about such entities has to proceed without the safety net of common sense.

Both of these factors - the absence of concrete instances of theoretical terms and the consequent lack of common sense knowledge about them - creates a strong tendency to *formalize* the scientific language game, to express the statements of a theory in a mathematical or logical calculus, so that the user of the theory need not rely on his intuitive understanding of its terms (Tarski, 1965). The rules of the calculus determine how the terms are to be used within the theory. Theoreticians can then focus their attention on the rules. The rules can, of course, be written down and so form concrete targets for agreements and disagreements. They can be discussed and constructively analyzed and investigated without universal agreement about exactly what intuitive content should be assigned to the terms they govern. Formalization also facilitates thinking. Mathematics, for instance, can be seen as a huge set of abstract inference-patterns which have been thoroughly investigated and found to be valid. Expressing a scientific theory in mathematical symbols allows the scientist to draw upon and apply those inference patterns while reason about his subject matter. A formal calculus supports thinking where common sense is powerless. In short, although formalization makes a subject matter incomprehensible to anyone who has no mastery of the particular formalism used, it nevertheless facilitates communication and collaborative thinking among the scientists who work with the theory.

Scientific theories, like ice-cream, come in different flavours. Throughout the diverse fields of scientific activity, different modes of knowing are employed, as they best fit the subject matter at hand (Toulmin, 1972). In some fields, it is common to build a ***model*** of the system being studied. If we are studying a system S, then a model of S is some other system M which is like S in some way. The question of exactly how M should be "like" S to be a good model is a question of some complexity (Hesse, 1970; Shanin, 1972). A model M is not merely "similar" to the system S. If this were the case, it would be unclear how the model could be useful in investigating

S. What could we discover by studying M other than the very similarity between M and S which made us select M as a model of S in the first place? The answer is that a good model has to be a generative system, i. e., a system which produces behavior, results, products, etc., of some sort. By investigating how the model M works - what results it generates under different circumstances, how its parts and properties interact - and by hypothesizing that the system S works in the same way, we can improve our understanding of S.

The term "theory" is used in both a narrow and a wide sense. In the narrow sense, a theory consists of a collection of "laws" - principles, statements, assertions - and, possibly, proofs of theorems, all of which have logical relationships to each other. A theory in this sense is basically a description, a text, even when expressed in some esoteric symbolism. A theory in the narrow sense contrasts with a model. A model is more like a replica; it represents the system of interest by duplicating its essential structure. It consists of components which have functional relationships to each other. A model is basically a picture, even when drawn in some non-visual medium like a programming language.

In the wider sense, a theory consists of both a set of principles and a model. The traditional astronomical theory of the solar system is an example. On the one hand, there is a system of equations which describes how the heavenly bodies act on each other and from which one can derive descriptions of how they move. This is a theory in the narrow sense. On the other hand, there is a geometric model of the solar system, consisting of spheres travelling along ellipses. The set of laws and the model together make up the astronomical theory in the wide sense. I will rely on context to make clear whether the term "theory" is used in the narrow or in the wide sense.

Only rarely does a scientific theory spring fully developed from its creator's head. More typically, a core idea is born, judged promising, and then worked out in an example. If the example is convincing, various developments follow. As the idea - fledging theory, as it were - is

applied to new situations, other, auxiliary ideas are combined with it, already established theories are compared to it, alternative formulations of it appear, efforts are undertaken to generalize it, etc.

In popular accounts of science, the creation of new theories is often portrayed as the main activity in research. In philosophical discussions, empirical testing of theories tends to occupy center stage. But the bulk of a scientist's time is spent neither in inventing theories nor in testing them, but in applying them. A theory is a complicated, abstract description. It is often far from self-evident how it should be used in order to construct an explanation for a particular fact or event. For instance, it took physicists more than two centuries to work out the many applications of Newton's theory of gravitation.

The activity of developing and applying a theory is very complex, both intellectually and sociologically (Toulmin, 1972). It is regulated by a paradigm (Kuhn, 1970) or a ***research programme*** (Lakatos, 1978), an "ideology" or point of view, loosely defined, which is more or less shared among those who work with the theory. A research programme consists of beliefs about a theory - a rationale for why the theory is interesting - and canons and conventions for working with that theory. It answers questions like the following: What is the core idea in the theory and which are the auxiliary ideas? What kinds of phenomena, and what kinds of observations of those phenomena, does the theory address? What kind of explanation does the theory provide? What counts as having explained something? When is an explanation successful? How can the theory be extended or improved? What form does progress take? What is the ultimate goal of the programme? What would count as "complete success"? Different programmes provide different answers to such questions for a particular field of research. In this essay, I am treating computer simulation as a research programme in the sense of Lakatos.

The Computer Simulation Programme

As noted in the introduction, computer simulation models are currently used by two different, although related, research programmes within cognitive psychology. The older programme conceptualizes cognitive processes as operations on symbol structures and builds simulations based on what we know about knowledge, while the more recent connectionist programme builds simulations based on what we know about the brain. The connectionist programme is reviewed in a separate section.

Psychologists sometimes use the term "simulation" to refer to computer programs which are not intended as theoretical models. In some experimental studies, subjects are observed while interacting with a simulation, e. g., a simulation of some complicated piece of machinery. In this case, the simulation is a model of the environment; it is a stimulus, or a technique for presenting a complicated stimulus. Also, some psychological theories are expressed in mathematical equations which are so complicated that they cannot be solved with the help of algebraic manipulations. In such cases, a computer program can be used to crunch out numerical solutions to those equations; such a program is often called a simulation[2]. In this essay, I will not discuss either of these two types of simulations.

The present section is divided into four parts: a statement of the rationale of the simulation programme, a summary of the workmode implied by that rationale, a select - and educationally biased - review of research inspired by it, and, finally, a critique.

The rationale

The core conception behind the use of computer programs as models of human cognitive processes consists of three main ideas. The first idea is that knowledge-relevant processes such as perception, remembering, and thinking operate upon internal symbol-structures, i. e., sense-

[2]See Loftus (1985) for critical comments on some uses of this kind of simulation.

impressions, memories, concepts, thoughts, ideas, images, etc., which symbolize or *represent* the external world. This stance - so self-evident to common sense as to be almost invisible for lack of contrasting alternatives - was rejected by behaviorist psychology, but was re-asserted throughout psychology and related disciplines during the fifties and early sixties. The notion that internal representations constitute the grist for the mental mill lies at the heart of modern cognitive psychology.

The second intellectual cornerstone of computer simulation is more technical and its historical origin more circumscribed. It is the insight that a computer has no intrinsic relation to numbers. The binary digits "1" and "0" which featured so prominently in early popularizations of computers can equally well be interpreted as standing for "true" and "false", or even "up" and "down", as for the numbers one and zero. Furthermore, groups of binary digits can be used as a code for any other symbol one might care to use, such as, say, the word "table", the implication-sign of formal logic, or the chemical symbol for oxygen. The computer is a general machine for operating upon symbols, *any* symbols. This conclusion appeared among a small number of individuals during the fifties (John McCarthy, Marvin Minsky, Allen Newell, Herbert Simon, among others). Their names are well-known to us because their insight made research into Artificial Intelligence (A. I.) possible.

The third and final step in the path to computer simulation as a psychological technique was taken by a team of three individuals and was first published in a joint paper entitled "Elements of a theory of human problem solving" (Newell, Shaw, & Simon, 1958). The paper combines the two previously described ideas: If human cognition operates on internal representations, and if computers can manipulate arbitrary symbol-structures, then computer programs are candidate models of thought processes.[3]

---

[3]The reader might want to compare the brief account given here with the more extensive discussion in the historical addendum in Newell & Simon (1972, pp. 873-889), and with Gardner (1985).

We should notice and dismiss two common misconceptions. First, a computer model is not a model of the ***brain***, but a model of the ***mind*** (or the cognitive aspect of mind). Computer simulation models are not interpreted in physiological terms like "brain cells", "electric potentials", etc., but in psychological terms like "knowledge", "attention", "memories", "thoughts", "beliefs", "concepts", "intentions", "inferences", "insights", "skills", etc.[4] Second, the organization of a general-purpose digital computer is determined by very different considerations than that of modelling the human mind. Consequently, the model does not reside in the physical machine, but in the program. In short, the claim of the simulation programme is not that computers are like brains, but that the process of executing a stored program shares something essential with the process of thinking.

To clarify what that shared essence is, we must take a brief look at the nature of computer programs. Programming languages are not assertive or descriptive but ***exhortational*** languages. They are designed to facilitate clear and unambiguous statement of procedures. A program is constructed by combining simple instructions into more complex ones, according to the rules of the particular programming language. Each instruction exhorts the machine to manipulate some symbol-structure in some way, e. g., moving it, changing it, deleting it, etc. Complex instructions are then combined into even more complex instructions, etc. A computer program is thus like a recipe, describing the procedure for, say, making apple pie, or a repair manual for an automobile, except that the objects being manipulated are symbols rather than apples or engines. From the psychological point of view, programs are like ***cognitive strategies***: problem solving strategies, attentional strategies, rehearsal strategies, memorization strategies, retrieval strategies, etc., of which modern cognitive psychology is full. If we have a (sufficiently specific) hypothesis about what strategy a person is applying to some cognitive task, we can specify that strategy in a programming language. Since the computer can read and execute such a specification (as long as

[4]One feature of the new connectionist programme is that its practitioners do try to discuss their models in physiological terms.

the conventions of the particular programming language have been followed precisely), loading the program into the machine enables the machine to apply that strategy. Program execution and thinking are similar in that both processes consist in the application of a strategy to a task (Simon, 1981).

A number of consequences follow from this basic conception. For instance, if a computer and a person are applying the same strategy to the same task, they should succeed and fail on the same problems. They should go through the same intermediate stages or produce the same partial results. They should make similar errors under similar circumstances. They should find the same problems easy or difficult, respectively. These expectations form the basis for the empirical evaluation of simulation models, a topic which will be discussed further below.

The simulation idea generates in-principle or generic answers to the basic questions of cognitive psychology. Every theoretical programme in psychology, I claim, has to have in-principle answers to the following four basic questions: Why is one task more or less difficult than another (for a given individual)? Why is one and the same problem easy for one individual and difficult for another? What is the nature of cognitive change? What aspects of cognition are invariant across tasks, across individuals, and over time? The generic answers provided by the computer simulation programme can be summarized as follows:

1. Inter-item differences. The generic explanation for task difficulty provided by the simulation programme is that one task requires more computations than the other. The amount of computation needed is, in turn, a function of a number of factors, such as the structure of the mental representation of the task and the need to avoid short-term memory load (see, e. g., Kotovsky, Hayes, & Simon, 1985).

2. Inter-individual differences. According to the simulationist's view, if two individuals perform differently on a task, then they are applying different methods to that task. The method can be analyzed into the representation of the task, the cognitive operations applied to it, and the strategy which controls the application of those operations, any of which may be responsible for a particular difference (see, e.g., Just & Carpenter, 1985; Newell & Simon, 1972; Ohlsson, 1984a).

3. Change. According to the simulation programme, cognitive change is the same as

revision of the knowledge-base in the head. Improvement in skills comes about through processes, so-called learning mechanisms, which revise - re-program, as it were - cognitive strategies (see Anderson, 1981a, and Klahr, Langley, & Neches, in press, for collections of articles about the simulation of skill learning). Improvement in declarative knowledge comes about through storage processes which extend the data-base in long-term memory. Forgetting is conceptualized as either decay of stored information or as interference between several pieces of information at the time of retrieval (see, e. g., Kolodner, 1984).

4. Generality. According to the computer simulation programme, it is the "mental programming language", or, more precisely, the structure of the information processing system in the head which remains invariant across individuals, across tasks, and over time (Newell, 1972; 1973a; 1973b).

Current research in cognitive psychology consists, to a large extent, of attempts to work out the details of these generic answers for particular experiments, situations, or phenomena.

The methodological advantages of the simulation technique can be organized into two groups, each group representing an issue with several different facets. First, there is the issue of *completeness*. If a simulation program solves a task correctly, then we know that the set of mechanisms in it are sufficient for the task. In the opinion of the simulationist, a psychological explanation must pass this sufficiency test before it can be seriously considered, i. e., before it is meaningful to compare it to empirical observations. If a subject solves an experimental task correctly, then the cognitive machinery in his head is thereby proven to be sufficient for the task, and, consequently, *any* explanatory machinery which is not sufficient is known to be incomplete, even before it is confronted with data. Sufficiency is a logical criterion for the adequacy of an explanation, similar to the criterion of consistency with respect to theories in mathematical physics. (An inconsistent set of equations is not an adequate physical theory.) Logical (as opposed to empirical) criteria for the adequacy of theories have not played an important role in the history of psychology, so it is not surprising that the sufficiency argument seldom is carried out or mentioned in actual simulation work, although it figured prominently in the argumentation of the programme founders (Newell & Simon, 1965).

In another facet of completeness, the simulation programme emphasizes that human beings are complete agents. Strictly speaking, there are no "perceptual tasks", "memory tasks", "problem solving tasks", etc. Perception, memory, and thinking are involved in every cognitive performance. It follows that there cannot be a theory of problem solving, or of memory, or of perception, because such a theory could not be tested against data; only complete systems, which have the entire range of capabilities (albeit, perhaps, in simplified form), and which shows explicitly the interactions between them, only such theories, the simulationist claims, can be meaningfully compared to empirical data.

Next, there is the issue of ***application***, which also has several different facets. One essential use of a theory or a model is to derive predictions from it. In the case of a simulation model, predictions can be derived by running the model on the computer. This is a convenient and inter-subjectively valid way of deriving predictions to be tested in experiments. We can also apply the model to situations which we have ***not*** observed (and perhaps could not observe), and so predict what the simulated person would have done in those situations. Other kinds of thought experiments are also possible. For instance, Ohlsson (1980) deleted randomly selected portions of the knowledge-base of a simulation program. Runs with the so mutilated program showed that the smaller knowledge-base did not cause the program to have fewer correct answers on a standard set of problems, but that the program had to work more in order to reach those answers. One can imagine running a simulation model under varying assumptions of short-term memory capacity, under assumption of perfect long-term memory recall, or under any number of extreme conditions, in order to investigate the implications of the principles embodied in it.

In summary, the computer simulation programme starts with the idea that a human being who is trying to cope with an intellectual task is applying a strategy to that task. If the strategy can be embedded in a computer program, then the process of executing that program on a computer can serve as a model of the mental process of the human while solving that task. This

research programme provides generic answers to the problems of why there are performance differences across problems, across individuals, and over time, as well as a principled conception of what remains invariant across those dimensions. The main methodological advantages of the simulation technique are that the medium forces the theoretician to produce complete explanations, and that the model, once constructed, can be applied to derive predictions and perform thought experiments.

The simulationist's workmode

Suppose that we have observed a particular person - a school child, a chess player, a doctor - while struggling with some intellectual task - a subtraction problem, an end-game, a tricky diagnosis - and that it is our goal to explain the observed performance. What shape does this activity take?[5]

The simulationist will begin with task analysis, i. e., with an effort to acquire as thorough an understanding of the task as possible. This might involve solving the task himself, perhaps repeatedly, and reflect on his own solution to it. It might involve reading and studying the knowledge domain to which the task belongs, perhaps even interviewing experts in the relevant field. What knowledge is necessary in order to handle the task? What knowledge is useful? What would be a useful organization of that knowledge? What different methods apply to the task? Eventually, the theoretician should be able to write a computer program for the task, execute it, and verify that it does, in fact, solve the task.

Two important aspects of simulation research are already evident. First, the simulationist must concern himself with the specific *content* of knowledge. Vague notions of "amount" of knowledge do not suffice for the writing of a program; the knowledge to be included in a program must be specified precisely. Second, the simulationist is likely to attend to *procedural* knowledge.

---

[5]The reader might want to compare the description given here with Kieras (1985) and Young (1985).

The methods, strategies, heuristics, tricks, rules of thumb, shortcuts, tactics, recipes, routines, and procedures of the relevant domain are of central interest because they constitute so many ideas about how a program might operate. Again, vague notions of, say, "cognitive style" do not suffice; the methods to be included in a program must be fully specified. In short, the writing of programs forces the simulationist to consider strategic knowledge, to consider the content of that knowledge, and to consider it in almost insufferable detail.

Once a running program has been constructed, the next task is to add psychological considerations. How do humans typically solve the relevant task? What changes should be made in the program to make it behave more like a person? When a first approximation has been achieved, it should be run on the same task as the human subject and a detailed trace of the program's behavior recorded.

Next, there follows the comparison of the trace of the simulation model with the behavior of the simulated person. The success of the model is measured by how closely its trace reproduces the observed performance. For instance, if a program is to be a good model of, say, a certain chess-player's skill, then it should certainly make the same chess-moves as that player when presented with the same chess-board configurations. Notice, in particular, that if the player makes mistakes, i. e., overlooks a strong move, then we would expect a good simulation model of him to overlook that same move. The basic evaluation principle is that if the subject gives answer A (correct or incorrect) to problem P, then the simulation model should produce answer A to problem P.

Of course, one can write many different computer programs which produce a particular answer to a particular problem. The construction of a simulation model needs to be informed and constrained by a wider data base than a single response. One approach to constraining model construction is to increase the temporal density of information about the particular performance

to be simulated by using detailed recordings of human performance, such as think-aloud protocols, recordings of eye-movements, and video-tapes of action sequences. The purpose is to achieve such a rich description of the performance so that differences between the computer trace and the human performance can easily be detected, and so that success on the part of the simulation is unlikely to come about accidentally. The evaluation principle used here is that the greater the amount of detail the simulation can reproduce or explain, the better it is (Newell & Simon, 1972).

A second approach to constraining a simulation model is to simulate performance on a range of related tasks. The purpose is, again, to achieve such a rich description of performance that success in modelling is unlikely unless the model captures the essence of the underlying mental processes. The evaluation principle used here is that the wider the range of performances a simulation can explain (at a given level of detail), the better it is (Kosslyn, 1980). Once the simulation model is confronted with data and the inevitable discrepancies noted, the simulationist enters into a sequence of successive revisions and re-comparisons, until the model has reached the desired level of accuracy.

In summary, the shape of the research activity generated by the simulationist idea is as follows: (a) Select a task to study. (b) Analyze and study the task until it is thoroughly understood. Write a computer program which can solve the task. (c) Collect observations from one or more human subjects solving the same task. (d) Revise the program in the way suggested by the data. (e) Run the program and compare its trace with the behavioral record. (f) Ponder what the implications are of the mis-matches between the computer trace and the behavioral record, and try to revise the model so that it reproduces the behavioral record better. The criteria for a successful explanation are that the trace of the program reproduces the human performance in some detail. The longer and more complicated the performance and the closer the detail in which it is reproduced, the better the simulation. Furthermore, if the model can

reproduce an entire set of related performances, then the larger the set, the better the simulation. The issues involved in the empirical evaluation of simulation models have been discussed by Kieras (1984, 1985), by Newell & Simon (1972), especially in chapters 5 and 6, and by Tuggle & Barron (1980), among others.

The activity of constructing an explanation within the simulation programme is a synthetic, rather than analytic, activity: an explanation is achieved by inventing and putting together a system, a device, which produces the event to be explained. Explaining is akin to ***designing***. A model is built by making choices on a number of design issues (Moore & Newell, 1974), e. g., what knowledge the subject had, what method the subject was using, what aspects of the task he attended to, what is the organization of the subject's knowledge, how are procedures represented in the subject's head, how is procedural knowledge used, what memory stores are there and what capacity limitations, if any, are they subject to, etc. Each simulation model represents a particular combination of choices on these and similar issues. A combination of choices constitutes a specification of the simulation program; computer implementation of that specification is analogous to the building of a physical device from engineering specifications.

The function of empirical observations in this research programme is different from its role in traditional experimental psychology. In a traditional study, the purpose of observing a person's performance is to ascertain the effect of some stimulus variable on his responses. In the simulation programme, the function of empirical investigations is, broadly speaking, to describe human performance in as much detail as possible, in order to constrain model building.

As the cycle of choosing tasks, observing human subjects, programming and testing models, continues, the simulation programme claims, we successively build up knowledge of what processes are necessary to accurately mimic human performance. In this research programme, knowledge takes the form of ***know-how***. Advances in knowledge will show up in the form of

increased ability to construct successful simulations. In the end, the simulationist claims, we will know enough about human cognition so that given only minimal information about an individual, we can quickly and easily put together a simulation model for his performance on some task, a model which will simulate that performance to some chosen degree of accuracy. Analogous to a physicist who can predict the trajectory of a particle, given its mass, its speed, and the direction in which it is moving, a simulationist will be able to predict a "cognitive trajectory", a solution path, given certain initial parameters (e. g., the subject's representation of the task, his previous knowledge, etc.).

Selected Review of Simulation Studies

The computer simulation idea originated within the study of problem solving. Over the three decades since the original Newell, Shaw, and Simon (1958) paper, the range of problem solving tasks used in simulation studies has grown continuously. Some of these tasks are puzzles with little or no pedagogical interest (e. g. Atwood, Masson, & Polson, 1980; Karat, 1982; Ohlsson, 1984b; Reed & Simon, 1976), but the last decade has witnessed a strong and growing trend to study class-room tasks, especially problems taken from mathematics and natural science. Detailed analyses of problem solving strategies is a hallmark of this strand of simulation research.

A second strand of simulation research is inspired less by ideas about thinking than about memory and language. Human long-term memory is a huge mass of diverse, irregular, vague, and incomplete descriptive knowledge which is typically extended, maintained, and used through acts of verbal communication, such as reading text or answering questions, and which is prone to forgetting. The use of semantic networks of various types to describe this knowledge base and its associated storage and retrieval strategies, and the willingness to implement large systems, are hallmarks of this second strand of simulation research.

In the following, I will briefly review four groups of simulation studies, three of which belong to the first strand described above, while the fourth group belongs to the second strand. In

a fifth and final miscellaneous group, I will mention some studies which do not fit neatly into either of the two main strands. The reader is warned against treating the present summary as a comprehensive overview. It is focussed on simulations with pedagogical relevance.

Cognitive development. The computer simulation community turned its attention to cognitive development early on. Several studies carried out in the seventies simulated childrens' behavior on classical Piagetian tasks, such as seriation of weight (Baylor and Gascone, 1974; Baylor, Gascone, Lemoyne, & Pothier, 1973) and of length (Young, 1976), classification and sorting (Klahr & Wallace, 1970), class inclusion (Klahr and Wallace, 1972), conservation of quantity (Klahr & Wallace, 1973), transitive inference (Klahr & Wallace, 1976), object constancy (Prazdny, 1980), etc.. The choice of Piagetian tasks was natural, given the amount of interest in Piaget's work in the sixties and seventies. The approach of these studies was to make detailed observations of childrens' behavior at the various developmental stages, and then write a separate simulation for each stage. The sequence of simulation models as a whole constitutes a representation, a slide-show, as it were, of the growth of mastery of the task. By comparing the successive models in the sequence, one can describe the developmental advances in theoretical, rather than behavioral, terms. One major conclusion from this mini-tradition is that the differences in the details of cognitive strategies ***within*** so-called developmental stages are at least as dramatic as the differences ***across*** such stages.

Although simulation studies of Piagetian tasks and phenomena are still carried out (see, e. g., Nason, 1986), current simulations of cognitive development generally show less of a Piagetian influence both in the choice of experimental tasks and in the theoretical questions asked. Issues about problem solving strategies, task encoding, and working memory capacity are explored in the context of tasks ranging from arithmetic to the Tower of Hanoi puzzle, and to the debugging of LOGO programs (Carver & Klahr, in press; Klahr, 1985; Klahr & Robinson, 1981; Siegler, 1986; Siegler & Shrager, 1984).

Most simulations in the field of cognitive development are not, strictly speaking, simulations of development, but simulations of childrens' (as opposed to adults') performance. They simulate performance at one or more stages of intellectual development, but they do not yet simulate the passage from one stage to another. The exception to this rule is the work by Wallace, Klahr, and Bluff (in press), which attempts to design an integrated system capable, in principle, of reproducing human cognitive development.

Performance models for educationally relevant problem solving tasks. Recently, simulationists have turned their attention to tasks which appear in the curriculum. Mathematics provides a store of well-defined tasks at many different levels of complexity. Arithmetic (Ashcraft, 1983; Briars & Larkin, 1984; Fletcher, 1985; Greeno, Riley, & Gelman, 1984; Hiebert & Wearne, in press; Klahr, 1973; Neches, 1981, 1982; Resnick & Neches, 1984; Riley & Greeno, 1980, Riley, Greeno, & Heller, 1983; Siegler, 1986; Siegler & Shrager, 1984; VanLehn, 1983), algebra, (Greeno, Magone, Rabinowitz, Ranney, Strauch, & Vitolo, 1985; Neves, 1978; Paige & Simon, 1965; Sleeman, 1982, 1984) and geometry (Anderson, 1981b, 1982; Anderson, 1983a; Anderson, Greeno, Kline & Neves, 1981; Greeno, 1976; Greeno, 1978a;) have all been targets for simulation efforts in recent years. Since the notion of an "error" comes naturally to the discussion of these domains, the design of computer programs which do mathematics incorrectly has become a major research activity. Many studies try to account for errors by showing how they would arise from the execution of some mathematical procedure which deviates from the correct procedure in some precisely specified (and sometimes minor) way. A major conclusion of this research is that errors cannot be thought of as resulting from ***missing*** knowledge only; incorrect mathematical strategies in students are not just incomplete, they also contain "mal-rules", i. e., components which are not part of the correct skill (Anderson, Boyle, & Yost, 1985; Brown & Burton, 1978; Burton, 1982; Evertz, 1982; Ohlsson & Langley, in press; Sleeman, 1982, 1984; Young & O'Shea, 1981).

A second area of the curriculum which has been mined for simulation targets is natural

science, especially physics problem solving (Bhaskar & Simon, 1977; Larkin, 1981; Larkin, McDermott, Simon, & Simon, 1980a, 1980b; Larkin, Reif, Carbonell & Gugliotta, in press). One approach used in this research is to simulate both a (successful) novice and an expert problem solver, and then characterize the differences between their solutions. One conclusion from such research is that expert problem solvers in physics employ mental representations of problems which are specifically ***physical***, while novices attend too much, or too soon, to the mathematics of the problem.

Human performance on non-academic but educationally important tasks like typing, electronic trouble shooting, and the operation of complicated machines of various kinds (including computers) have also been investigated in simulation studies (Anzai, 1984; Card, Moran, & Newell, 1983; Kieras & Polson, 1985; Rumelhart & Norman, 1982; White & Frederiksen, 1984). Such studies are about to mature into a new kind of applied psychology with strong implications for workplace design and personnel training.

The acquisition of cognitive skills. Given a sequence of models each of which simulates a person's performance at different developmental stages or at different points in time, the next question to ask is what mechanisms affect the transition from one to the other. In the computer simulation approach, this question is, of course, approached by incorporating learning mechanisms into simulation models. Computer learning is still a rather novel phenomenon even within the field of Artificial Intelligence, and running simulation programs which simulate how humans learn cognitive skills are few and far between. The volumes by Anderson (1981a) and by Klahr, Langley, & Neches (in press) cover a significant portion of the relevant research[6] .

A small group of simulation studies have tried to model human learning in artificial task domains. Anzai and Simon (1979) simulated one person's successive mastery of the Tower of

[6] See Mitchell, Carbonell, & Michalski (1986) for an overview of relevant Artificial Intelligence research.

Hanoi puzzle. Langley (1983a, 1985) has applied a theory of discrimination learning to several puzzle tasks. In a similar effort, Ohlsson (in press-a) simulated the acquisition of strategies for several puzzle tasks, as well as the transfer of training from one group of problems to another within a very simple task domain. Feigenbaum and Simon (1984) have proposed a general theory of learning to recognize patterns, such as nonsense syllables. Rosenbloom and Newell (1986) has suggested a general scheme for human learning called "chunking" which has been used to simulate performance on laboratory reaction-time tasks, with excellent fit to human data. Ohlsson (in press-b) has designed a learning mechanism which simulates successive changes in the mental representation of a simple verbal reasoning task, guided by general world knowledge.

However, there is also a handful of studies simulating the acquisition of cognitive skills relevant to classroom tasks. John Anderson and coworkers have simulated learning to construct proofs in elementary plane geometry (Anderson, 1982; Anderson, 1983a; Anderson, Greeno, Kline, & Neves, 1981) and to write simple computer programs (Anderson, 1986; Anderson, Farrell, & Sauers, 1984). Kurt VanLehn's simulation of the acquisition of the standard algorithm for multi-column subtraction takes into account an entire lesson-sequence, taking a lesson to be a set of solved examples and letting the simulation program learn from them (VanLehn, 1983). Neches (1981, 1982) has simulated the discovery of the so-called MIN-strategy for simple addition[7]. Neves (1978) has simulated the learning of algebra rules from solved examples.

Learning, retaining, and recalling declarative knowledge. The simulations discussed so far are all concerned with procedural knowledge, i.e., knowledge about how to act in order to reach certain goals. A second main strand of simulation research is mainly interested in declarative knowledge, i.e., knowledge about what is true of the world. The paradigmatic task in this tradition is to read a text (perhaps only a single sentence), and then, after some period of time,

[7]Empirical research has established that children can discover the MIN-strategy on their own (Groen & Resnick, 1977).

answer one or more questions about it. A simulation model for this task has to contain assumptions about how long-term memory is organized, about how new knowledge is added to it, and about how information is retrieved from it. It is desirable, albeit difficult and laborious, to provide such a simulation with a natural language front-end which models how people understand language and which allows the model to acquire its knowledge by reading actual texts. This type of simulation model grew out of the idea of semantic networks (Quillian, 1968, 1969) and flourished in the seventies (Anderson, 1976; Anderson & Bower, 1973; Collins & Quillian, 1969; Frijda, 1972; Kintsch, 1974; Miller, 1981; Norman, Rumelhart, & the LNR group, 1975; Rumelhart, Lindsay, & Norman, 1972; Schank, 1972, 1975; Schank & Abelson, 1977; Schank & Colby, 1973). The reviews by Frijda (1972) and by Chang (1986) provide further discussion about this strand of simulation research.

One of the main results of research with simulation models based on natural language interfaces and semantic networks is that systematic and reasonable strategies for storage and retrieval of declarative information can produce the kinds of confusions, transformations, and retrieval failures which are typical of human performance (Kolodner, 1983a, b, 1984). Another lesson from this type of research is that acquisition and utilization of declarative knowledge is always context dependent. For instance, a question is answered differently depending on when and where it is asked, by whom, etc. The complexity of declarative knowledge acquisition may be responsible for what seems to be a recent decline in this strand of simulation research: current simulations of memory retrieval tend to shy away from implementing the actual content of human long-term memory in favor of more abstract, and therefore more computationally tractable, models (Hintzmann, 1984; Raajmaker & Shiffrin, 1981; Walker & Kintsch, 1985; Wiegerswa, 1982.).

Simulations of long-term memory have, in principle, tremendous pedagogical relevance. In school subjects like history, geography, and social studies, "learning" is, to a large extent,

"studying", i. e., acquiring mostly descriptive knowledge through reading or listening. The simulation models mentioned here contain explicit hypotheses about the acquisition, retaining, forgetting, and recalling of declarative knowledge. Clearly, such hypotheses must have implications for how declarative knowledge should be presented to facilitate learning. However, in spite of their obvious pedagogical relevance, simulation models within this strand have so far not figured frequently in educational applications. One exception to this rule is the work by Kolodner (1983c) in which she applies considerations of long-term memory to the investigation of how expertise in psychiatric diagnosis is acquired.

Miscellaneous. There are other simulation studies which deserve mention as well, although they do not fit neatly into any of the above categories. Of particular interest from a pedagogical point of view are simulations of reading (Just & Carpenter, 1984; Just & Thibadeau, 1984; Kieras, 1982, 1983; Thibadeau, Just, & Carpenter, 1982). The efforts to simulate skill acquisition has spawned a handful of studies trying to simulate initial language acquisition (Anderson, 1983b, chap. 7; Langley, 1982; Wolff, 1980). A particularly fascinating application of the natural language/semantic network approach to computer simulation is the effort by Colby (1973, 1975, 1981) to simulate neurotic and paranoiac thought processes. Another type of model tries to simulate human performance on various visualization tasks. This mini-tradition was initiated by Baylor (1973), has been intensively developed by Kosslyn (1980, 1985), and continues to generate new ideas about the mind's eye (Just & Carpenter, 1985; Morgan, 1983). Finally, the reader might be interested to know that even such a difficult-to-grasp process as daydreaming is now being implemented on computers (Mueller & Dyer, 1985a, b).

Critique of the Simulation Programme

Three decades worth of experience with the simulation programme should be enough to form at least a tentative evaluation of it. A scholarly appraisal of the programme cannot be attempted here. What follows is the view of one participant and practitioner.

It may at first glance seem as if the simulation programme has been spectacularly successful in furthering our understanding of cognition. Today, the entire field of cognitive psychology speaks the language of symbolic computation; the psychological journals are full of mental representations being encoded, stored, retrieved, and processed in all manner of ways. However, a critique of the simulation idea must distinguish between the information processing conception of mental processes, on the one hand, and the simulation idea proper, on the other. The value of seeing cognitive processes as symbolic computations is not in doubt; at least, it will not be questioned here.[8] The question is whether the simulation *technique*, i. e., the explanatory procedure which leads to implemented, runnable computer programs, provides additional advantages over and above the advantages of adopting the information processing *point of view*.

When the question is posed in this way, the case for computer simulation is weak indeed. It is difficult to think of a single example of a surprising, important, or interesting *discovery* which has been made by running a simulation model. It is equally difficult to think of an example of some dramatic *prediction* which has been made by a simulation program, a prediction which was subsequently born out by data. Nor is it easy to state the *principles* presumably established by, say, the simulation studies mentioned in the last subsection. As far as this author knows, no great *theoretical synthesis* has been achieved by taking two simulation programs and combining them into a new model with more predictive power than either of the parts. Finally, there seems to be little or no growth of *consensus* about what mechanisms one must postulate in order to have a reasonable simulation of a human thought process.

These discouraging observations demand an explanation. We need to scrutinize the difficulties and disadvantages with the simulation programme, distinguishing as we do so between the following three questions:

1. Have the practitioners carried out the programme carefully?

[8] Connectionism does, to some extent, pose that question.

If the practitioners have been careless or neglectful or simply misunderstood the simulation technique, the lack of results would not speak against the programme itself.

2. To what extent have the practitioners been limited by their tools?
   Computer simulation can only proceed where Artificial Intelligence has cleared the way. If current programming techniques are inadequate to capture some aspect of human knowledge, the failure to simulate that aspect cannot be chalked up as failure of the simulation programme. The limitations of the programme itself has to be distinguished from the limitations of our current resources for carrying it out.[9]

3. What are the *inherent* weaknesses, inconsistencies, and difficulties with the simulation programme?

Keeping these three questions in mind, we turn to a brief discussion of the problems associated with, respectively, convenience code, disembodied processes, effort, brittleness, empirical validation, restricted public access to models, and informal statement of theories.

Convenience code. Any simulation program will contain code which helps make the program executable but which does not represent any psychological hypotheses. For instance, a simulation of chess playing might contain a routine which finds all legal moves in a chess board configuration. The psychologist may not have any hypotheses about *how* the simulated chess player finds the legal moves; the theory might be all about choosing between moves, once found. Thus, the psychologist writes any convenient piece of code which will accomplish the move finding task, and lets his theoretically motivated code select a move by operating on the output from the former.

The problem is that the motivated code and the convenience code interact in producing the behavior of the program. We have no guarantee that the performance of the program depends in essential ways on the theoretically motivated parts of the code. For instance, suppose the move finding routine is faulty, so that it ignores a certain class of weak but legal moves. The

[9]The situation here is similar to that which holds in physics. Physicists are dependent upon mathematical results in order to develop their theories - with the result that physicists sometimes make substantial contributions to mathematics. Similarly, psychologists with computer simulation on their agenda sometimes have to dig into front-line A. I. research problems in order to develop the tools they need.

theoretically motivated move-choosing routine will never choose one of those weak moves, not because it is intelligent enough to avoid them, but because it never gets to consider them. The program "predicts" that the chess player will never make that particular type of error. If this "prediction" is born out by observations, it will be taken as support for the theory behind the program, although it actually depends on the design of the convenience code. The definition of a "legal move" in chess is sharp enough so that the particular accident described in this illustrative example is unlikely to happen. But other, equally disastrous interactions between motivated code and convenience code may be opaque enough to go unnoticed by the programmer. This problem has been discussed by Frijda (1967), Kieras (1985), and by Neches (1982).

The architecture-program distinction was proposed by Allen Newell (1972; 1973a; 1973b) as a solution to this problem. He argues that simulationists should create psychologically motivated programming languages in which to write their models. The very structure of such a language - its architecture - should be based on psychological hypotheses, so that the programs written in it need not contain any convenience code. ***All*** code in a simulation program should, in principle, admit of psychological interpretation.

The two most ambitious efforts to date to create psychologically motivated programming languages are the ACT architecture by John Anderson (1976, 1983b) and the Soar architecture by Laird, Rosenbloom, & Newell (1986). Approaching the problem in a slightly different way, the PRISM programming language is an attempt to provide a flexible architecture, in which certain aspects, e.g., the number of working memories, can easily be changed to accommodate changes in theoretical assumptions (Langley, 1983b). However, the call by Allen Newell to distinguish between architecture and program and to write simulation programs in such a way that their every feature admits of psychological interpretation, this call has, to a large extent, gone unheeded.

Disembodied processes. A human being is a complete agent, capable of perception, thinking, remembering, motor action, and learning. Clearly, all these capabilities enter into the solution to any cognitive task; observable behavior is a result of the interaction between them. Any explanation of behavior must describe the role of each process in that interaction.

Psychologists, however, tend to be more interested in one of these capabilities than in the others. When they turn to computer simulation, they, naturally enough, want to simulate the capability they are most interested in. The result is often a simulation of what we might think of as a disembodied process. For instance, simulations of problem solving often ignore the question of how the subject's thinking interacts with his long-term memory; what is the problem solver reminded of while thinking about the problem? Similarly, most simulations of long-term memory can answer questions through retrieval of information, but not through problem solving. Simulations of either problem solving and question answering usually ignore issues of learning. Yet another type of incompleteness which is virtually universal is to ignore the perceptual-motor interaction with the environment. Simulation programs regularly assume that the entire task description is present and accessible in working memory, finessing the issues of attention allocation and of physical manipulation of task materials.

Taking a purist stance, I conclude that such simulations are useless, because we already know that the observable behavior of a person comes about through an interaction between perception, memory, thinking, motor action, and learning. The predictions of a program which represents one of these processes remain indeterminate until its interactions with the other processes have been specified. Every simulation should have a complete set of mental processes, i. e., mechanisms for remembering, thinking, and learning, a (simulated) environment, and capabilities for perceptual as well as a motor interaction with that environment. The simulation programme calls for simulations of agents, not of disembodied processes.

Effort. Writing a serious simulation model is no small enterprise. Several months have to be spent analyzing the task, many more in implementing a program which can solve it, and, quite likely, many more still in revising it until it actually mimics a human subject. If the target task is at all complex, considerable A. I. know-how might be needed to make a runnable program.

This fact has unfortunate consequences for how the simulation programme is carried out. First, there is, quite naturally, a reluctance to explore alternative explanations of a phenomenon. If getting one simulation model up and running is laborious, the idea of doing it twice is not attractive. To the best of my knowledge, no researcher to date has implemented two radically different simulation models for the same set of observations and then compared them. (Comparisons between two different theories with respect to the same phenomenon is, of course, standard procedure in other domains of science.) Second, there is a certain reluctance to make use of other people's models. I know of no example of a researcher constructing an explanation of his own data by re-implementing somebody else's simulation model from the published accounts of it, and applying it to those data. (Again, using other researchers' theories in constructing explanations of new data is standard procedure in other domains of science.)

The problem of effort has no radical solution. We can only hope that more powerful A. I. programming tools and accumulation of know-how within the simulation community will eventually shrink to reasonable proportions the amount of work that has to go into each new simulation model.

Brittleness. One of the major difficulties with A. I. programs at the current time is that a program which works well on one task often fails dramatically on conceptually similar tasks. Since computer simulation models are limited to what can be done with current A. I. programming techniques, simulations are brittle as well. This has the unfortunate consequence of preventing thought experiments and empirical explorations of programs. Making

changes in a program and then running it can be a frustrating experience, because it may turn out that the changes render the program impossible to execute. Exploration of the generality of a model can be discouraging, because there may be precious little in the model which transfers to a new task domain. In short, applying the program to other problems than the ones it was designed for may turn out to be impossible.

However, brittleness is not a methodological problem, but a theoretical one; it is a consequence of our ignorance of the nature of intelligence. If we understood how human minds avoid being brittle, we would also understand how to write non-brittle simulations. Understanding intelligence is, of course, our goal. Brittleness and its associated difficulties are not properties of the simulation methodology, but indications of our progress, or lack of progress, with respect to our research agenda.

Empirical validation. Perhaps the most commonly noted problem with simulation models is that they are severely underdetermined by the data they are applied to. For any model, it seems, one can think of an alternative model which accounts equally well for the data. The large number of assumptions in a simulation model and the complexity of their interactions - so the argument goes - make it very difficult, if not downright impossible, to test a model stringently against empirical observations.[10]

The first thing to notice about this argument is that the possibility of alternative explanations for a phenomenon is not a feature of the simulation programme, but an aspect of all scientific research. Since other branches of science have coped successfully with this difficulty, the argument does not have any force unless it can be shown that the problem of choosing between rival explanations is inherently more severe with simulation models than with other kinds of theories. No such demonstration has yet been provided.

[10] John Anderson has argued that this is necessarily so; see Anderson (1979).

The second thing to notice is that current practice with respect to empirical validation of simulation models is unsatisfactory. The use of temporally dense behavioral records like think-aloud protocols (Ericsson & Simon, 1984), eye-movement recordings (e.g., Just & Carpenter, 1985), and video-tapes of actions (e.g., Klahr & Robinson, 1981) is still uncommon. Even worse, when those methods are used, the sequential information they contain is usually destroyed by categorization or aggregation procedures of various kinds. Since the main difference between two different models for the same task is the sequence of intermediate results that they produce, ignoring sequential information is fatal. Furthermore, the methodological insight that it is a more challenging task to simulate a single performance in detail than to simulate the global characteristics of average performance (Newell & Simon, 1972) has largely gone unnoticed. Cognitive psychology has adopted the ***theoretical*** tool of computer simulation, but it has, by and large, not adopted the associated ***empirical*** methodology of temporally dense recordings of single performances.

It appears, then, that simulation models are more underconstrained by data than they need to be. The empirical validation of simulation models is a problem, not because of difficulties peculiar to the simulation programme, but because of the nature of scientific research in general and because of the reluctance of psychologists to adopt the empirical methodology that was invented for that very purpose.

Restricted public access to models. A computer program is not easy to read, its code is usually too voluminous to publish, and it is sometimes difficult to get a program developed on one computer to run on another. Consequently, simulation models are essentially private. They are described in scientific publications, but they are seldom used in a serious way by others than their creators. To some critics, this feature of simulation contradicts the public nature of scientific knowledge (Frijda, 1967; Neches, 1982).

This criticism is correct in its statement of the facts and it indicates a serious difficulty with the simulation programme, but I think the critics have misinterpreted the nature of that difficulty. To clarify the issue, we have to make use of the distinction between theories and models. A model does not, in and of itself, represent knowledge. A model is a particular, in philosophical terminology; in fact, being itself a thing, it is every bit as particular as the thing it models. A simulation model does not "say" anything at all about human cognition; it does not make pronouncements or assertions; like all other objects, it just exists. The activities of designing, implementing, and applying a model can help produce knowledge, but that knowledge resides in the head of the model-builder, not in the model. It is certainly essential to make that knowledge publicly accessible; it is more doubtful if the same is true of the model itself.

Consider the double-helix constructed by Watson and Crick as a model of the DNA molecule (Watson, 1968). In no sense was the ***model itself*** - the rickety construction of differently colored pieces of metal - made public. They did not ship it to other laboratories, nor would such a road show have served any purpose. Concrete, physical models used in scientific research are usually private. The important thing is that the corresponding ***theory*** is given public expression.

In summary, the criticism that simulation programs are private tools is true, but it makes the wrong point. It treats programs as if they were theories rather than models. A model may remain private, as long as the theory it embodies does not. As we shall see below, the crux of the matter is exactly how to make the theory public.

Informal statement of theories. A programming language is a formal notation in which one can describe a concrete model of information processing. However, there is no corresponding formal language for expressing the theory behind that model. Therefore, simulationists are forced to use the only resource available, namely natural language. The psychological principles which

guide the construction of the formal models are themselves informally stated.[11]

This has several consequences. First, it is unclear by what criterion a model is judged to instantiate a particular theory. If two researchers disagree whether a certain model embodies such-and-such a theory, how would one decide the issue in a disciplined way? (This difficulty raises the same question as the problem of convenience code discussed above: how do we know that the runs of a particular simulation program are relevant to the theory we are interested in?) Second, since the theory behind a simulation model is stated informally, we cannot expect the simulation technique to produce the advantages of formalized theory construction that were mentioned in the introduction.

Thus, the simulation programme is radically incomplete. It provides a methodology for the construction of concrete models of information processing *mechanisms*, but it does not provide a tool for the statement and analytical treatment of information processing *principles*. What is needed is a high-level, formal, so-called specification language in which psychological principles can be stated and deductively analyzed and from which specific simulation models can be generated in a systematic way.[12]

Hagert (1986) has proposed a solution to this problem, based on logic programming techniques. In his method, a program specification (i. e., a theory) is stated in formal logic. Procedures (i. e., models) which satisfy that specification can be generated by deduction. The procedures are directly executable in the Prolog programming environment, so they can be run and empirically validated. The deductive derivation of a procedure (model) from a specification (theory) ensures that (a) the procedure satisfies the specification, and (b) that all the assumptions

---

[11] The reader may want to confirm this by, e. g., looking at Chapter 14 of Newell and Simon (1972), where they state their theory of problem solving, or the statement by Anderson (1983b), pp. 17-38, of the ACT theory.

[12] The new approach of connectionism has such a two-layered theoretical methodology.

behind the procedure have been stated explicitly. Hagert (1986) has applied this method to the simulation of a simple short-term memory task. It remains to be seen whether the method is useful in the construction of more complicated simulation models.

Summary. Although the simulation programme, as distinct from the information-processing point of view, has ad a dramatic effect on psychological research, it has not been a spectacularly successful enterprise. This is due, in part, to the fact that its practitioners have not carried out the program carefully, and, in part, to limitations on the available programming tools. However, the main weakness is that the simulation programme is radically incomplete. It offers a technique for building concrete models, but no methodology for stating in a precise form what we learn from building those models. We can only express the theories we arrive at in informal language. Consequently, research within the simulation programme does not benefit from the advantages of formal theorizing.

Connectionism: The Return of Mathematics

Some cognitive psychologists are currently engaged in working out a research programme which might supersede computer simulation as described on this essay. This new research programme is commonly known as ***connectionism*** (Feldman & Ballard, 1982). The reader is referred to Hinton and Anderson (1981) and McClelland and Rumelhart (in press a, b) for representative works, as well as to the special issue of the journal *Cognitive Science* (Volume 9, Issue 1, 1985).

The connectionist approach is based on the idea that we should use what we know about the brain, as well as what we know about behavior, when we try to simulate the mind. Perhaps the most basic observation one can make about brain activity is that many parts of the brain are active at any one time. This leads to a conception of distributed information processing, where many processing units operate in parallel. According to this view, information processing in the head is being done by "small" units, each of which performs a very simple processing task, like

computing the sum of the strengths of its inputs, and then transmits a signal to a certain set of other such units. There are many units, each of which has a large number of connections to other units. The units all act locally, on their own inputs only, so they can operate in parallel. The network of units is characterized by rules which regulate the connections between units, the type of computation being performed by each unit, and the kind of signal sent from one unit to an other. A learning theory in this programme consists of a set of rules for how the units should change their characteristics, e. g., which signal to send, as a result of their own activity. An information processing system built along these lines is rather different in character from the symbol manipulation systems considered so far in this essay.

How would connectionism answer the four questions I used to introduce the traditional simulation programme?

1. Inter-item differences. A connectionist theory is well suited to describe task differences which are due to interference effects and confusion; indeed, explaining such phenomena is its paradigmatic application. When it comes to large differences in performance, it is less obvious how the connectionist type of theory applies. Why, for instance, should one isomorph of the Tower of Hanoi problem take 16 times longer than another to solve (Kotovsky, Hayes, & Simon, 1985)? Why is the "envelope" version of Wason's so-called selection task so much easier to solve then the "card" version (Evans, 1984)?

2. Inter-individual differences. There are no particular conceptual tools within the connectionist scheme to explain differences between individuals.

3. Change. Connectionism regards learning, or change over time in general, as a central characteristic of cognition. A connectionist learning theory consists essentially of a rule for how to change one or more of the characteristics of a processing unit, as a function of the inputs it has received in the past.

4. Generality. The invariants of the mind in the connectionist programme are the network of units, the function which maps the input signals onto an output signal, and the rule for how to change that function.

In short, the connectionist programme has principled answers to the questions about change and about invariants, but does not handle differences between items or between individuals well.

Methodologically, the connectionist approach has two strong advantages over the

"traditional" computer simulation approach. First, it combines formal *model-building* with formal *theorizing*, in a way which computer simulation as described earlier does not do. Connectionist models can be implemented on a computer and simulation runs, just like other simulation models. But due to the uniformity of the processing units and the simplicity of the computations they perform, the behavior of an entire connectionist network can be described and analyzed with mathematical tools. Thus, there is a formal way of stating the general properties of connectionist computer models.[13]

The second methodological advantage of connectionism is that it can use both behavioral and neurophysiological data to constrain model-building. As noted in a a previous section, one of the problems with computer simulation models is that they are obviously underdetermined by the empirical data they are intended to explain. By letting their simulation models be models of brain activity, the connectionists make it possible to apply what is known about the brain to constrain those models.

The practitioners of the connectionist programme characterize their object of study as the "micro-structure of cognition". The tasks they have studied are mostly perceptual-motor and short-term memory tasks. This accords with a long tradition in psychology of studying tasks for which one's theoretical approach seems to have the highest face validity: perceptual-motor processes obviously happen in parallel and in a distributed fashion. The educational implications of a deeper understanding of human performance on these tasks are not obvious. Application of connectionist modelling techniques to tasks like proof finding in geometry and learning the history of World War II is yet to come.

The most disturbing aspect of connectionism is the fact that it, once again, takes the

---

[13]The technique of doing a mathematical description of a computer simulation programme has also been used in connection with more traditional simulation programs which use so-called spread of activation as one of its memory retrieval processes; see Anderson (1984) and Anderson & Pirolli (1984).

knowledge out of the theory of cognition. The point of the connectionist approach - and its most fascinating aspect - is that it describes a type of system in which knowledge has no location, in which knowledge is spread out among a large number of units, and in which knowledge only exists in so far as the entire set of units reacts as a whole in particular ways to particular situations. Consequently, connectionism does not provide any tools for the analysis of knowledge. We cannot, for instance, command a computer to print out a connectionist model and expect to find its knowledge in the print-out; the only thing that can be printed is a very long list of connections with their associated strengths, as incomprehensible as the brain itself. Thus, connectionism takes an anti-representational stance. There are no mental representations in the head, no symbols which refer to the external world, nor any cognitive strategies. To the extent that the post-war advances in our understanding of cognition are due to thinking of mental processes in terms of operations on internal representations, under the guidance of heuristic knowledge to that extent those advances are lost again in the connectionist programme.

Impact of Simulation on Educational Research and Practice

Three decades of research have produced some knowledge about how to design simulation models of cognitive processes. What impact will this ability have on educational research and practice?[14] The question turns out to have several layers, a *methodological* layer, a *theoretical* layer, and a *technological* layer. Each layer will be discussed below. The final conclusion is that the main consequence of bringing cognitive psychology into contact with education is that practice gets a chance to influence research, rather than the other way around.

The computer simulation technique has two methodological consequences which contribute to its educational relevance. The first is that tasks used in psychological experiments are now chosen on the basis of inherent interest. Experimental tasks used to be chosen on practical

[14]The broader question of what the implications are for education of modern cognitive psychology in general (rather than just the simulation technique *per se*) has been considered by Gagne (1985), Frederiksen (1984) and Ohlsson (1983); see also the collection of articles by Tuma & Reif (1980).

grounds: They were convenient to administer, involved cheap materials that could be systematically varied, demanded responses that could easily be registered, etc. Presumably, the practitioners of traditional experimental psychology believed that since they were interested in general laws of behavior, they did not have to pay attention to the particular properties of any one task. A psychologist interested in computer simulation, however, must attend to the particular content of the mental representations and processes involved in solving the experimental task. Consequently, selecting an interesting task has become a central step in the posing of a cognitive research problem. Since educational relevance is one ground for regarding a task as interesting, a significant proportion of current cognitive research concerns tasks which appear in the curriculum or in training programs of one form or another. Cognitive psychologists now do their research on the very same tasks which teachers try to teach.

A second educationally interesting consequence of the simulation technique is the new willingness of psychologists to study single subjects. The experimental psychology of, say, 1920-1960 dealt almost exclusively with averages over groups of subjects.[15] But educators benefit very little from knowing that people *in general* have such-and-such characteristics. Educational practice depends upon the ability to react properly to particular individuals. The simulation technique allows single individuals to be studied and modelled, which makes it relevant for questions like "What does this student know?'", "Why did this student make that error?", "What kind of instruction does this student need?", etc.

Taking the relationship between physics and chemistry, on the one hand, and engineering science, on the other, as our analogy, we would expect that better theories about human cognition in general and about human learning in particular would lead to changes in educational practice. According to this view, the influence of the simulation technique on education would be indirect:

[15]Before that time, analyses of single subjects were more common. Thus, in this respect, modern cognitive psychology is returning to an earlier methodological stance; see Dukes (1965).

as a research tool, it will contribute to better theories of human cognition, and these, in turn, will contribute to the improvement of instruction.

Examples of such influences at the theoretical layer are not hard to find. For instance, consider the notion of an arithmetic error. Due to the research by Brown and Burton (1978) and Burton (1982), the vague idea of an arithmetic algorithm not being properly "fixed" in memory has given way to the notion of a "buggy procedure". The latter obviously derives from the conception of cognitive skills as analogous to computer programs. Computer implementations of buggy algorithms constitute concrete, detailed, and precise descriptions which can be used to predict the behavior of a student who is the victim of a particular error. The increased understanding of cognitive errors produced by such simulations facilitate a precise discussion of how to diagnose errors and what remedial instruction to provide.

As a second example, consider the question why cognitive skills are difficult to learn. Anderson's (1982, 1986) theory of knowledge compilation show and expect him to apply to it immediately. rule, say, and expect him to apply it immediately. The translation of a verbally stated piece of advice into an executable procedure which interacts correctly with previously existing procedures is seen to be a complex process, once we try to spell it out in a computer simulation. Similarly, VanLehn's (1983) theory of procedure induction explains why a student might arrive at the wrong algorithm as a result of studying correctly solved examples: The induction of the correct procedure from the examples presupposes knowledge of what the relevant properties of the examples are, knowledge which students do not always possess.

In short, the simulation technique does contribute to an improved understanding of the cognition of classroom tasks, with implications for instruction. However, such knowledge transfer is slow and uncertain; teachers do not always pay attention to research results, and researchers do not always communicate them well. Also, scientific breakthroughs do not occur every day. This

form of interaction between research and pedagogy is weak and the effects are likely to be small, at least in the short run.

However, science does not only - and perhaps not primarily - influence human activity through the spread of new knowledge, but through the technologies supported by that knowledge. The main impact of simulation on pedagogy, I suggest, will be mediated by the *teaching devices* of various kinds which are soon to be appear in classrooms. Current know-how supports two types of devices which have computer models as essential ingredients: systems for automatic cognitive diagnosis (Anderson, Boyle, & Yost, 1985; Attisha & Yazdani, 1984; Burton, 1982; Johnson, 1985; Johnson & Soloway, 1984; Marshall, 1980, 1981; Ohlsson & Langley, in press; Reiser, Anderson, & Farrell, 1985; Sleeman, 1982, 1984) and intelligent tutoring systems (Anderson, Boyle, Corbett, & Lewis, 1986; Anderson, Boyle, & Reiser, 1985; Beakley & Haden, 1985; Sleeman & Brown, 1982; Wenger, 1985). Such devices are beginning to appear in schools and their presence there will grow more frequent over time.

Like any other type of tool, these tools will shape the activity in which they are used. To clarify the nature of this effect, let us consider, as an illustrative example of a socially influential technology, that ever-present communication tool, the telephone. The effect of the telephone on our communication habits is a result of the functionality of the device, not of its scientific underpinnings; the telephone does not influence people by making them think of electricity as they dial, nor is wide-spread understanding of, or even belief in, the theory of electricity a necessary condition for the telephone to have an impact. Similarly, the new teaching devices will shape instruction in ways which depend primarily on the activities that the devices allow. For instance, intelligent tutoring systems can monitor practice in a constructive way, freeing the teacher from the task of correcting essays and worksheets; while systems for automatic cognitive diagnosis enable a teacher to have a detailed picture of the knowledge state of each student in her class. We do not yet know what effects such devices will have on teaching. The point argued here

is that the effects, whatever they turn out to be, will be caused by the functions of these tools - to monitor practice, to perform cognitive diagnosis - rather than by the content of the psychological theories which motivate their design. Furthermore, the appreciation of those functions does not presuppose a general understanding of the underlying theories. Hence, the slowness of knowledge transfer between the research community and the practitioners does not limit the impact of such devices on education.

The discussion so far has asked the question of how simulation might influence education. I now want to argue that the teaching devices made possible by computer simulation open up a channel through which educational practice might come to decide the fate of psychological theories and set the agenda for psychological research. Consider the fact that tutoring systems, once placed in the classroom, constitute the most efficient tool for collecting precise and ecologically valid observations about human cognition that has ever been devised. Unprecedented amounts of empirical information can be collected with very little effort. We can foresee at least the following effects on research:

1. New learning phenomena will be discovered. It would be truly surprising if data collection on the scale that is now becoming possible did not reveal previously unnoticed regularities in human learning.

2. The probability that educational research will be side-tracked by artifactual experimental effects will decrease. A proposed regularity or phenomena can be checked against a vast, ecologically valid data-base.

3. Interest in laboratory studies will decline. What educational researcher - indeed, what cognitive psychologist - does not prefer to collect his data in the classroom? Why not use ecologically relevant data if they are available?

4. The sheer amount of empirical material reported in each journal article is likely to increase. In the long run, the standards of journal reviewers and editors will be affected.

5. The type of studies performed are likely to change. On the one hand, data collection in the classroom favors descriptive studies, because many aspects of classroom activity cannot be varied systematically for the sake of an investigation.

   On the other hand, intelligent tutoring systems can be programmed to teach with

different teaching strategies, and they are guaranteed to execute a particular teaching strategy faithfully. This enables well-controlled experimental evaluations of teaching strategies, as well as experimental tests of theoretically derived predications about the effects of different teaching strategies on the behavior of the learner. Perhaps "teaching strategy" will become *the* major independent variable in future psychology experiments.

6. If most psychological observations are collected in the classroom, the only way for a theory to be impressive is to handle those observations. Classroom behavior, rather than laboratory behavior, will become the main testing ground for theories of cognition. (See Anderson, Boyle, Corbett, & Lewis, 1986, for a similar opinion.).

At the present time, these effects of computerized teaching devices on cognitive research are just beginning to be noticeable.[18]

In summary, the interactions between the simulation technique and educational research and practice are complex. The simulation technique forces researchers to consider the content of their subjects' knowledge and provides a format for theorizing about individual subjects. Both of these features contribute to the pedagogical relevance of cognitive psychology. We can expect a traditional science-to-practice knowledge transfer to occur to the extent that simulation models contribute to the improvement of psychological theories with pedagogical relevance. But simulation know-how also facilitates the construction of computerized teaching devices. Like all other tools, these devices will shape the activities in which they are being used through their functionality, rather than through their scientific rationale. Finally, these devices constitute a channel through which the behavior of students in real learning situations become available to researchers, with dramatic effects on how cognitive research is carried out.

[18] Interestingly, O'Dell & Dickerson, (1984), make a similar argument for the usefulness of systems for computerized psychotherapy such as the ELIZA program.

References

Anderson, J. R. (1976). ***Language memory and thought***. Hillsdale, NJ: Erlbaum.

Anderson, J. R. (1979). Arguments concerning representations for mental imagery. ***Psychological Review, 86***, 395 -406.

Anderson, J. R. (1980). ***Cognitive psychology and its implications***. San Francisco: Freeman.

Anderson, J. R. (1981a). ***Cognitive skills and their acquisition***. Hillsdale, NJ: Erlbaum.

Anderson, J. R. (1981b) Tuning of search of the problem space for geometry proofs. ***Proceedings of the Ninth International Joint Conference on Artificial Intelligence***, Vancouver, Canada.

Anderson, J. R. (1982). Acquisition of cognitive skill. ***Psychological Review, 89***, 369-406.

Anderson, J. R. (1983a). Acquisition of proof skills in geometry. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), ***Machine learning. An artificial intelligence approach***. Palo Alto: Tioga Publishing Co.

Anderson, J. R. (1983b). ***The architecture of cognition***. Cambridge, MA: Harvard University Press.

Anderson, J. R. (1984). Spreading activation. In J. R. Anderson & S. M. Kosslyn (Eds.), ***Tutorials in learning and memory***. San Francisco: W. H. Freeman.

Anderson, J. R. (1986). Knowledge compilation: The general learning mechanism. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), ***Machine Learning. An artificial intelligence approach*** (Vol. II). Los Altos, CA: Kaufmann.

Anderson, J. R., Boyle, C. F., Corbett, A., & Lewis, M. (1986). Cognitive modelling and intelligent tutoring. *Technical Report*. Carnegie-Mellon University, Psychology Department.

Anderson, J. R., Boyle, D. F., & Reiser, B. J. (1985). Intelligent tutoring systems. *Science, 228*, pp. 456-462.

Anderson, J. R., Boyle, D. F., & Yost, G. (1985). The geometry tutor. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*.

Anderson, J. R., & Bower, G. H. (1973). *Human associative memory*. Washington, DC: V.H. Winston.

Anderson, J. R., Farrell, R., & Sauers, R. (1984). Learning to program in LISP. *Cognitive Science, 8*, 87-129.

Anderson, J. R., Greeno, J. G., Kline, P. J., & Neves, D. M. (1981). Acquisition of problem-solving skill. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Erlbaum.

Anderson, J. R., & Pirolli, P. L. (1984). Spread of activation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 791-798.

Anzai, Y., & Simon, H. A. (1979). The theory of learning by doing. *Psychological Review, 86*, 124-140.

Anzai, Y. (1984). Cognitive control of real-time event-driven systems. *Cognitive Science, 8*, 221-254.

Ashcraft, M. H. (1983). Simulating network retrieval of arithmetic facts. (Tech. Rep. #10). Pittsburgh: University of Pittsburgh, Learning Research and Development Center.

Attisha, M., & Yazdani, M. (1984). An expert system for diagnosing children's multiplication errors. *Instructional Science, 13*, 79-92.

Atwood, M. E., Masson, M. E. J., & Polson, P. G. (1980). Further explorations with a process model for water jug problems. *Memory and Cognition, 8*, 189-192.

Bhaskar, R., & Simon, H. A. (1977). Problem solving in semantically rich domains: An example from engineering thermodynamics. *Cognitive Science, 1*, 193-215.

Baylor, G. W. (1973). Modelling the mind's eye. In A. Elithorn & D. Jones (Eds.), *Artificial and human thinking*. Amsterdam: Elsevier.

Baylor, G. W., & Gascone, J. (1974). An information processing theory of aspects of the development of weight seriation in children. *Cognitive Psychology, 6*, 1-40.

Baylor, G. W., Gascone, J., Lemogne, G., & Pothier, N. (1973). An information processing model of some seriation tasks. *Canadian Psychologist, 14*, 167-196.

Beakley, G. E., & Haden, C. R. (Eds.). (1985). *Computer-aided processes in instruction and research*. New York: Academic Press.

Briars, D. J., & Larkin, J. H. (1984). An integrated model of skill in solving elementary word problems. *Cognition and Instruction, 1*, 246-296.

Brown, J. S., & Burton, R. R. (1978). Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science, 2*, 155-192.

Burton, R. (1982). Diagnosing bugs in a simple procedural skill. In D. Sleeman & J. S. Brown (Eds.), *Intelligent tutoring systems* (pp. 157-183). London: Academic Press.

Card, S. K., Moran, T. P., & Newell, A. (1983). *The psychology of human-computer interaction*. Hillsdale, NJ: Erlbaum.

Carver, S. M., & Klahr, D. (in press). Assessing children's LOGO debugging skills with a formal model. *Journal of Educational Computing Research*.

Chang, T. M. (1986). Semantic memory: Facts and models. *Psychological Bulletin, 99*(2), 199-220.

Colby, K. M. (1973). Simulations of belief systems. In R. S. Schank & K. M. Colby (Eds.), *Computer models of thought and language*. San Francisco: Freeman.

Colby, K. M. (1975). *Artificial paranoia: A computer simulation of paranoid processes*. New York: Pergamon.

Colby, K. M. (1981). Modeling a paranoid mind. *Behavioral and Brain Sciences, 4*, 515-560.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior, 8*, 240-247.

Dukes, N. F. (1968). N = 1. *Psychological Bulletin, 64*, 74-79.

Ericsson, K. A., & Simon, H. A. (1984). *Protocol analysis. Verbal reports as data*. Cambridge, MA: The MIT Press.

Evans, J. S. (1984). Heuristic and analytical processes in reasoning. *British Journal of Psychology, 75*(4), 451-468.

Evertz, R. (1982). ***A production system account of children's errors in fraction subtraction.*** (Tech. Rep. #28). Milton Keyes, U.K.: The Open University, CAL Research Group.

Feigenbaum, E., & Simon, H. A. (1984). EPAM-like models of recognition and learning. *Cognitive Science, 8*(4), 305-336.

Feldman, J. A., & Ballard, D. H. (1982). Connectionist models and their properties. *Cognitive Science, 2*, 205-254.

Fletcher, C. F. (1985). Understanding and solving arithmetic word problems: A computer simulation. ***Behavior Research Methods, Instruments, and Computers, 175***, 565-571.

Frederiksen, N. (1984). Implications of cognitive theory for instruction in problem solving. *Review of Educational Research, 54*(3), 363-407.

Frijda, N. H. (1967). Problems of computer simulation. *Behavioral Science, 12*, 59-67.

Frijda, N. H. (1972). Simulation of human long-term memory. *Psychological Bulletin, 77*, 1-31.

Gagne, E. (1985). ***The cognitive psychology of school learning.*** Boston: Little, Brown & Co.

Gardner, H. (1985). ***The mind's new science.*** New York: Basic Books.

Glaser, R. (1978). The contributions of B. F. Skinner to education and some counterinfluences. In P. Suppes (Ed.), ***Impact of research on education: Some case studies*** (pp. 199-265). Washington, DC: National Academy of Education.

Good, T. L., & Brophy, J. E. (1986). ***Educational psychology*** (3rd ed.). New York: Longman.

Greeno, J. G. (1976). Indefinite goals in well-structured problems. *Psychological Review, 83*,

479-491.

Greeno, J. G. (1978a). A study of problem solving. In R. Glaser (Ed.), ***Advances in instructional psychology*** (Vol. 1). Hillsdale, NJ: Erlbaum.

Greeno, J. G., Magone, M. E., Rabinowitz, M., Ranney, M., Strauch, C., & Vitolo, T. M. (1985). ***Investigations of a cognitive skill.*** (Tech. Rep. #27). Pittsburgh: University of Pittsburgh, Learning Research and Development Center.

Greeno, J. G., Riley, M. S., & Gelman, R. (1984). Conceptual competence and children's counting. ***Cognitive Psychology, 16,*** 94-143.

Groen, G. J. (1978). The theoretical ideas of Piaget and educational practice. In P. Suppes (Ed.), ***Impact of research on education: Some case studies*** (pp. 267-317). Washington, DC: National Academy of Education.

Groen, G. J., & Resnick, L. B. (1977) Can preschool children invent addition algorithms? ***Journal of Educational Psychology, 69,*** 645-652.

Hagert, G. (1986). ***Logic modelling of conceptual structures: Steps towards a computational theory of reasoning.*** Unpublished doctoral dissertation, University of Uppsala, Uppsala, Sweden.

Hesse, M. B. (1970). ***Models and analogies in science.*** Notre Dame: The University Press.

Hiebert, J., & Werne, D. (in press). A model of students' decimal computation procedures. ***Cognition and Instruction.***

Hinton, G. E., & Anderson, J. A. (Eds.). (1981). ***Parallel models of associative memory.***

Hillsdale, NJ: Erlbaum.

Hintzmann, D. L. (1984). MINERVA 2: A simulation model of human memory. *Behavior Research Methods, Instruments & Computers, 16*(2), 96-101.

Johnson, W. L. (1985, May). *Intention-based diagnosis of errors in novice programs*. (Tech. Report No. YALEU/CSD/RR #395). New Haven, CT: Yale University, Dept. of Computer Science.

Johnson, W. L., & Soloway, E. (1984). Intention-based diagnosis of programming errors. *Proceedings of the National Conference of the American Association for Artificial Intelligence* (pp. 162-168).

Just, M. A., Carpenter, P. A. (1984). Using eye fixations to study reading comprehension. In D. E. Kieras & M. A. Just (Eds.), *New methods in reading comprehensive research*. Hillsdale, NJ: Erlbaum.

Just, M. A., & Carpenter, P. A. (1985). Cognitive coordinate systems: Accounts of mental rotation and individual differences in spatial ability. *Psychological Review, 92*(2), 37-171.

Just, M. A., & Thibadeau, R. H. (1984). Developing a computer model of reading times. In D. E. Kieras & M. A. Just (Eds.), *New methods in reading comprehension research*. Hillsdale, NJ: Erlbaum.

Karat, J. (1982). A model of problem solving with incomplete constraint knowledge. *Cognitive Psychology, 14*, 538-559.

Kieras, D. E. (1982). A model of reader strategy for abstracting main ideas from simple technical phase. *Text, 2*, 47-82.

Kieras, D. E. (1983). A simulation model for the comprehension of technical prose. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 17). New York: Academic Press.

Kieras, D. E. (1984). A method of comparing a simulation model to reading time data. In D. Kieras and M. Just (Eds.), *New methods in reading comprehension research*. Hillsdale, NJ: Erlbaum.

Kieras, D. (1985). The why, when, and how of cognitive simulation: A Tutorial. *Behavior Research Methods, Instruments, and Computers, 17*(2), 279-285.

Kieras, D., & Polson, P. G. (1985). An approach to the formal analysis of user complexity. *International Journal of Man-Machine Studies, 22*(4), 365-394.

Kintsch, W. (1974). *The representation of meaning in memory*. Hillsdale, NJ: Erlbaum.

Klahr, D. (1973). A production system for counting, subitizing and adding. In W. G. Chase (Ed.), *Visual information processing*. (pp. 527-546). New York: Academic press.

Klahr, D. (1985). Solving problems with ambiguous subgoal ordering: Preschoolers' performance. *Child Development, 56*, 940-952.

Klahr, D., Langley, P., & Neches, R. (Eds.). (in press). *Production system models of learning and development*. Cambridge, MA: MIT Press.

Klahr, D., & Robinson, M. (1981). Formal assessment of problem-solving and planning processes in preschool children. *Cognitive Psychology, 13*, 113-148.

Klahr, D., & Wallace, J. G. (1970). An information processing analysis of some Piagetian experimental tasks. *Cognitive Psychology, 1*, 350-387.

Klahr, D., & Wallace, J. G. (1972). Class inclusion processes. In S. Farnham-Diggory (Ed.), *Information processing in children*. New York: Academic Press.

Klahr, D., & Wallace, J. G. (1973). The role of quantification operators in the development of conservation of quantity. *Cognitive Psychology, 4*, 301-327.

Klahr, D., & Wallace, J. G. (1976). *Cognitive development. An information processing view*. Hillsdale, NJ: Erlbaum.

Kolodner, J. L. (1983a). Maintaining organization in a dynamic long-term memory. *Cognitive Science, 7*, 243-280.

Kolodner, J. L. (1983b). Reconstructive memory: A computer model. *Cognitive Science, 7*, 281-328.

Kolodner, J. L. (1983c). Towards an understanding of the role of experience in the evolution from novice to expert. *International Journal of Man-Machine Studies, 19*, 497-518.

Kolodner, J. L. (1984). *Retrieval and organizational strategies in conceptual memory: A computer model*. Hillsdale, NJ: Erlbaum.

Kosslyn, S. M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.

Kosslyn, S. M. (1985, October). *Visual hemispheric specialization: A computational theory*. (Tech. Report #7). Cambridge: Harvard University.

Kotovsky, K., Hayes, J. R., & Simon, H. A. (1985). Why are some problems hard? Evidence from Tower of Hanoi. *Cognitive Psychology, 17*, 248-294.

Kuhn, T. S. (1970), *The structure of scientific revolutions* (2nd Ed.). Chicago, IL.

Laird, J. E., Rosenbloom, P. S., & Newell, A. (1986). *Universal subgoaling and chunking: The automatic generation and learning of goal hierarchies*. Norwell, MA: Kluwer.

Lakatos, I. (1978). *The methodology of scientific research programmes. Philosophical papers* (Vol. I). Cambridge, U.K.: Cambridge University Press.

Langley, P. (1982). Language acquisition through error recovery. *Cognition and Brain Theory, 5*, 211-255.

Langley, P. (1983a). Learning search strategies through discrimination. *International Journal of Man-Machine Studies, 18*, 513-541.

Langley, P. (1983b). Exploring the space of cognitive architectures. *Behavior Research Methods and Instrumentation, 15*, 289-299.

Langley, P. (1985). Learning to search: From weak methods to domain-specific heuristics. *Cognitive Science, 9*, 217-260.

Larkin, J. H. (1981). Enriching formal knowledge: A model for learning to solve textbook physics problems. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Erlbaum.

Larkin, J. H., McDermott, J., Simon, D. P., & Simon, H. A. (1980a). Expert and novice performances in solving physics problems. *Science, 208*, 1335-1342.

Larkin, J. H., McDermott, J., Simon, D. P., & Simon, H. A. (1980b). Models of competence in solving physics problems. *Cognitive Science, 4*, 317-345.

Larkin, J. H., Reif, F., Carbonell, J., & Gugliotta, A. (in press). FERMI: A flexible expert

reasoner with multi-domain inferencing. *Cognitive Science.*

Loftus, G. (1985). Johannes Kepler's computer simulation of the universe: Some remarks about theory in psychology. *Behavior Research Methods, Instruments, & Computers, 17,* 149-156.

Marshall, S. P. (1980). Procedural networks and production systems in adaptive diagnosis. *Instructional Science, 9,* 129-143.

Marshall, S. P. (1981). Sequential item selection: Optimal and heuristic policies. *Journal of Mathematical Psychology, 23,* 134-152.

McClelland, J. L., Rumelhart, D. E., & the PDP Research Group (Eds.). (in press,a). *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1). Cambridge, MA: Bradford Books/MIT Press.

McClelland, J. L., Rumelhart, D. E., & the PDP Research Group (Eds.). (in press,b). *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2). Cambridge, MA: Bradford Books/MIT Press.

Mitchell, T. M., Carbonell, J. G., & Michalski, R. S. (Eds.). (1986). *Machine learning: A guide to current research.* Norwell, MA: Kluwer.

Miller, J. R. (1981). Constructive processing of sentences: A simulation model of encoding and retrieval. *Journal of Verbal Learning and Verbal Behavior, 20*(1), 24-45.

Moore, J., & Newell, A. (1974) How can Merlin understand? In L. W. Gregg (Ed.) *Knowledge and cognition.* Potomac, MD: Erlbaum.

Morgan, M. J. (1983). Mental rotation: A computationally plausible account of transformation through intermediate steps. *Perception, 12*(1), 203-211.

Mueller, E. T., Dyer, M. G. (1985a, August). Daydreaming in humans and computer. *Proceedings of the Ninth Joint Conference on Artificial Intelligence* (Vol. 1) (pp. 278-280). Los Angeles, CA.

Mueller, E. T., Dyer, M. G. (1985b, August). Towards a computational theory of human daydreaming. *Proceedings of the Seventh Annual Conference of the Cognitive Science Society* (pp. 120-129), Irvine, CA.

Nason, R. (1986). *A production system analysis of young childrens' number development between the ages two to eight years.* Unpublished doctoral dissertation, Deakin University, Victoria, Australia.

Neches, R. T. (1981). A computational formalism for heuristic procedure modification. *Proceedings of the Seventh International Joint Conference on Artificial Intelligence.*

Neches, R. T. (1982). Simulation systems for cognitive psychology. *Behavior Research Methods & Instrumentation, 14,* 77-91.

Neves, D. M. (1978). A computer program that learns algebraic procedures by examining examples and working problems in a textbook. *Proceedings of the Second National Conference of the Canadian Society for Computational Studies of Intelligence.*

Newell, A. (1972). A theoretical exploration of mechanisms for coding the stimulus. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, DC: Winston.

Newell, A. (1973a). Production systems: Models of central structures. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press.

Newell, A. (1973b). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press.

Newell, A., Shaw, J. C., & Simon, H. A. (1958). *Psychological Review, 65*(3), 151-166.

Newell, A. & Simon, H. A. (1965). Programs as theories of higher mental processes. *Computers in biomedical research*.(Vol. 2). New York: Academic Press.

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Norman, D. A., Rumelhart, D. E. & The LNR Group. (1975). *Explorations in cognition*. San Francisco: Freeman.

O'Dell, J. W., & Dickson, J. (1984). ELIZA as a "therapeutic" tool. *Journal of Clinical Psychology, 40*(4), 942-945.

Ohlsson, S. (1980, January). *Competence and reasoning with common spatial concepts*. (Working Papers from the Cognitive Siminar No. 6). Stockholm, Sweden: University of Stockholm, Department of Psychology.

Ohlsson, S. (1983). The enaction theory of thinking and its educational implications. *Scandinavian Journal of Educational Research, 27*, 73-88.

Ohlsson, S. (1984a, June). Attentional heuristics in human thinking. *Proceedings of the Sixth Conference of the Cognitive Science Society*. Boulder, CO.

Ohlsson, S. (1984b). Induced strategy shifts in spatial reasoning. *Acta Psychologica, 57*, 47-67.

Ohlsson, S. (in press,a). Transfer of training in procedural learning: A matter of conjectures and refutations? In L. Bole (Ed.), *Computational models of learning*. Springer-Verlag.

Ohlsson, S. (in press,b). Truth vs. appropriateness: Relating declarative to procedural knowledge. In R. Neches, D. Klahr, & Langley (Eds.), *Self-modifying production systems: Models of learning and development*. Cambridge, MA: Bradford Books/The MIT Press.

Ohlsson, S., & Langley, P. (in press). Psychological evaluation of path hypothesis in cognitive diagnosis. In H. Mandl & A. Lesgold (Eds.), *Learning issues for intelligent tutoring*. New York: Springer Verlag.

Paige, J. M., & Simon, H. A. (1965). Cognitive processes in solving algebra word problems. In B. Kleinmuntz (Ed.), *Problem solving: Research, Method, and Theory*. New York: Wiley.

Prazdny, S. (1980). A computational study of a period of infant object-concept development. *Perception, 9*(2), 125-150.

Quillian, M. R. (1968). Semantic memory. In M. Minsky (Ed.), *Semantic information processing* (pp. 227-270). Cambridge, MA: MIT Press.

Quillian, M. R. (1969). The teachable language comprehender. *Communications of the Association for Computing Machinery, 13*, 459-476.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search in associative memory. *Psychological Review, 88*, 93-134.

Reed, S. K., & Simon, H. A. (1976). Modeling strategy shifts in a problem solving task. *Cognitive*

Psychology, 8, 86-97.

Reiser, B. J., Anderson, J. R., & Farrell, R. G. (1985). Dynamic student modelling in an intelligent tutor for LISP programming. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*. Los Angeles.

Resnick, L. B., & Neches, R. (1984). Factors affecting individual differences in learning ability. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (Vol. 2) (pp. 275-323). Hillsdale, NJ: Erlbaum.

Riley, M. S., & Greeno, J. G. (1980). *Details of programming a model of children's counting in ACTP*. (Tech. Rep. #6). Pittsburgh: University of Pittsburgh, Learning Research and Development Center.

Riley, M. S., Greeno, J. G., & Heller, J. I. (1983). Development of children's problem-solving ability in arithmetic. In H. P. Ginsburg (Ed.), *The development of mathematical thinking*. New York: Academic Press.

Rosenbloom, P. S., & Newell, A. (1986). The chunking of goal hierarchies: A generalized model of practice. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine Learning. An artificial intelligence approach* (Vol. II). Los Altos, CA: Kaufmann.

Rumelhart, D. E., Lindsay, D. H., & Norman, D. A. (1972). A process model for long-term memory. In E. Tulving & W. Donaldson (Eds.), *Organization and memory*. New York: Academic Press.

Rumelhart, D. E., & Norman, D. A. (1982). Simulating a skilled typist: A study of skilled cognitive-motor performance. *Cognitive Science, 6*, 1-36.

Schank, R. C. (1972). Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology, 3*, 552-631.

Schank. R. C. (1975). *Conceptual information processing*. Amsterdam: North-Holland.

Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding*. Hillsdale, NJ: Erlbaum.

Schank, R. C., & Colby, K. M. (Eds.). (1973). *Computer models of thought and language*. San Francisco: Freeman.

Shanin, T. (Ed.). (1972). *The rules of the game. Cross-disciplinary essays on models in scholarly thought*. London: Tavistock.

Siegler, R. S. (1986). Unities across domains in children's strategy choices. In M. Perlmutter (Ed.) *Minnesota Symposium of Child Psychology*, (Vol 19). Hillsdale, NJ: Erlbaum.

Siegler, R. S., & Shrager, J. (1984). Strategy choices in addition and subtraction: How do children know what to do? In C. Sophion (Ed.), *Origins of cognitive skills*. Hillsdale, NJ: Erlbaum.

Simon, H. A. (1981). Studying human intelligence by creating artificial intelligence. *American Scientist, 69*(3), 300-309.

Sleeman, D. (1982). Assessing aspects of competence in basic algebra. In D. Sleeman & J. S. Brown (Eds.), *Intelligent tutoring systems*. New York: Academic Press.

Sleeman, D. (1984). An attempt to understand students' understanding of basic algebra. *Cognitive Science, 8*(4), 387-412.

Sleeman, D., & Brown, J. S. (Eds.). (1982). ***Intelligent tutoring systems***. New York: Academic Press.

Tarski, A. (1965). ***Introduction to logic and to the methodology of deductive sciences***. New York: Galaxy.

Thibadeau, R., Just, M. A., & Carpenter, P. A. (1982). A model of the time cause and content of reading. ***Cognitive Science, 6***, 157-203.

Toulmin, S. (1972). ***Human understanding*** (Vol. 1). ***General introduction and Part 1***. Oxford: Clarenden Press.

Tuggle, F. D., & Barron, F. (1980). On the validation of descriptive models of decision making. ***Acta Psychologica, 45***, 197-210.

Tuma, D. T., & Roof, F. (Eds.). (1980). ***Problem solving and education: Issues in teaching and research***. Hillsdale, NJ: Erlbaum.

VanLehn, K. (1983). ***Felicity conditions for human skill acquisition: Validating an AI-based theory*** (Tech. Report No. CIS-21). Palo Alto, CA: Xerox PARC.

Walker, W. H., & Kintsch, W. (1985). Automatic and strategic aspects of knowledge retrieval. ***Cognitive Science, 9***, 261-283.

Wallace, I., Klahr, D., & Bluff, K. (in press). A self-modifying production system model of cognitive development. In D. Klahr, P. Langley, & R. Neches (Eds.), ***Models of learning and development***. Cambridge, MA: MIT Press

Watson, J. D. (1968). ***The double helix***. New York: Mentor Books.

Wenger, E. (1985). *AI and the communication of knowledge: An overview of intelligent tutoring systems*. (Tech. Report). Irvine, CA: University of California, Irvine Computational Intelligence Project, Dept. of Information and Computer Science.

White, B. Y., & Frederiksen, J. R. (1984, June). Modeling expertise in troubleshooting and reasoning about simple electric circuits. *Proceedings of the Sixth Annual Conference of the Cognitive Science Society* (pp. 337-343). Boulder, CO.

Wiegerswa, S. (1982). Sequential response bias in randomized response sequences: A computer simulation. *Acta Psychologica, 52*, 249-256.

Wolff, J. G. (1980). Language acquisition and the discovery of phrase structure. *Language and Speech, 23*(3), 253-269.

Young, R. M. (1976). *Seriation in children. An artificial intelligence analysis*. Basel, West Germany: Birkhauser Verlag.

Young, R. M., & O'Shea, T. (1981). Errors in children's subtraction. *Cognitive Science, 5*, 153-177.

Young, S. R. (1985). Programming simulations of cognitive processes: An example of building macrostructures. *Behavior Research Methods, Instruments and Computers, 17*, 286-293.

Author Notes

Preparation of this manuscript was supported in part by NSF grant MDR-8470339, and in part by an institutional grant for the OERI Center for the Study of Learning, OERI-G-86-0005, United States Office of Education. The opinions expression do not necessarily reflect the position of either NSF or OERI, and no official endorsement should be inferred. I thank the students in my graduate course on "Formal theories in psychology" for stimulating discussions about topics relevant for this essay.

Copyright © Stellan Ohlsson